Towards MRI-Based Autonomous Robotic US Acquisitions: A First Feasibility Study

Christoph Hennersperger*, Bernhard Fuerst*, Salvatore Virga*, Oliver Zettinig, Benjamin Frisch, Thomas Neff, and Nassir Navab

arXiv:1607.08371v1 [cs.RO] 28 Jul 2016

Abstract

Robotic ultrasound has the potential to assist and guide physicians during interventions. In this work, we present a set of methods and a workflow to enable autonomous MRI-guided ultrasound acquisitions. Our approach uses a structured-light 3D scanner for patient-to-robot and image-to-patient calibration, which in turn is used to plan 3D ultrasound trajectories. These MRI-based trajectories are followed autonomously by the robot and are further refined online using automatic MRI/US registration. Despite the low spatial resolution of structured light scanners, the initial planned acquisition path can be followed with an accuracy of 2.46 ± 0.96 mm. This leads to a good initialization of the MRI/US registration: the 3D-scan-based alignment for planning and acquisition shows an accuracy (distance between planned ultrasound and MRI) of 4.47 mm, and 0.97 mm after an online-update of the calibration based on a closed loop registration.

Index Terms

robotic ultrasound, autonomous acquisition, image-guidance, rgb-d camera, multi-modal registration

I. Introduction

ULTRASOUND (US) has become one of the standard medical imaging techniques and is widely used within diagnostic and interventional applications. During needle insertion for liver biopsy or ablation, ultrasound imaging is used today in order to guide the insertion process throughout the procedure. Based on pre-interventional X-ray Computed Tompgraphy (CT) or Magnetic Resonance Imaging (MRI) datasets, a needle is advanced carefully under US-guidance to reach a final target position. Thereby, it is important to maintain steady ultrasound views throughout the intervention, providing an overview of all essential structures in the abdomen (tumor, vessels, lung, liver tissue).

In general, clinical US is mostly based on 2D-images (except cardiac and gynaecological applications), requiring a manual navigation of the probe. The resulting high operator-variabilty of the manual guidance is not only challenging for the application described above, but impairs a wider clinical acceptance of ultrasound for the extraction of quantifiable parameters from these data [1]. 3D ultrasound can potentially overcome these limitations, and is performed either using native 3D probes or by tracking 2D images in space, using a tracking target attached to the ultrasound probe (tracked freehand 3D ultrasound) [2]. While systems using native 3D probes are still not widely available in clinical practice, tracked ultrasound is easily accessible and can also be interpolated with respect to a regular grid [3]. When comparing native and freehand 3D ultrasound, both techniques have their merits. Native 3D ultrasound allows, on the one hand, for live 3D volume acquisitions in real-time and thus for a direct analysis of volume changes over time; a property which is especially exploited for 3D echocardiography [4]. On the other hand, the systems are expensive and only allow for the imaging of comparably small volumes restricted by the probes' field of view. Freehand 3D ultrasound does not pose limitations with respect to volume sizes, anatomies and trajectories, but can be potentially distorted by varying pressure applied by the operator, or changing anatomy caused by breathing or cardiac pulsation.

While ultrasound acquisitions are mostly performed by physicians or medical staff, data collection for both 2D and 3D ultrasound might be improved by automating the acquisition process. This goes in hand with the demands of the healthcare system, showing that current freehand ultrasound scanning techniques suffer from high operator variability, hindering the overall clinical acceptance, especially of 3D applications [1], [5].

To this end, a fully automatic system to acquire 3D ultrasound datasets could avoid the high inter-operator variability and resource demands of current clinical systems. In view of liver ablation for example, such automatic ultrasound acquisitons could be performed periodically throughout the procedure to guide the insertion, while CT or MRI datasets will be available for an interventinal planning. By addressing limitations of today's ultrasound techniques using a robotic framework, this would not only open up the way for new applications in an interventional setting, but also for potential society-wide screening programs using non-invasive ultrasound imaging [1], [6]. In this view, autonomous acquisitions have the potential both to facilitate clinical acceptance of robotic imaging techniques by simplifying the workflow as well as to reduce the scanning time and the necessary amount of manual operator input.

*Authors contributed equally. Corresponding author e-mail: christoph.hennnersperger@tum.de — C. Hennersperger, S. Virga, O. Zettinig, B. Frisch, and N. Navab are with the Chair for Computer Aided Medical Procedures, Technische Universität München, Germany. T. Neff and S. Virga are further with KUKA Roboter GmbH, Augsburg, Germany, and B. Fuerst and N. Navab are further with Computer Aided Medical Procedures, Johns Hopkins University, Baltimore, USA.

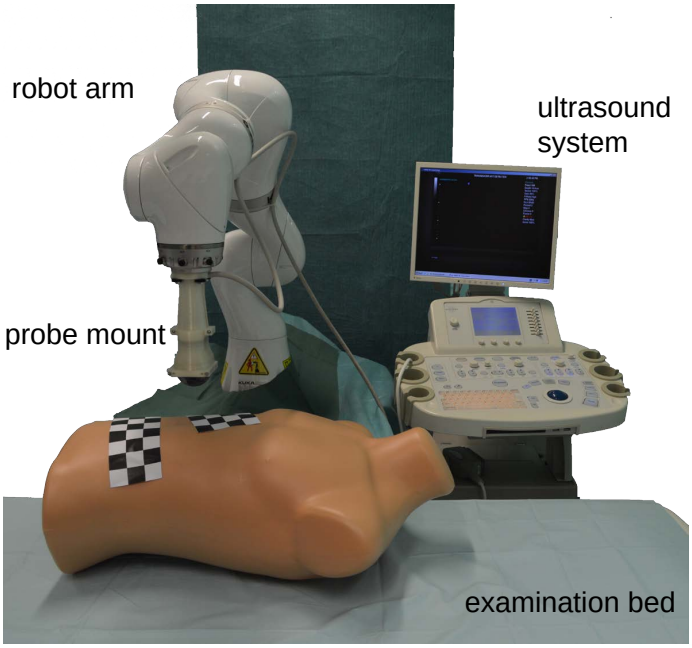

Fig. 1. In the presented system setup, a RGBD-camera is mounted on the ceiling, observing the scene with the patient lying on the examination bed. The robotic arm is approaching the patient from above and is equipped with a curvilinear ultrasound probe.

First attempts aiming at automatized ultrasound acquisitions used mechanical stepper motors, moving the ultrasound probe in a controlled fashion [7]. More recently, different robotic systems were developed in the context of clinical ultrasound, including applications to imaging, surgical interventions and needle guidance [8], [9]. While freehand ultrasound scanning enables a fast and dynamic acquisition and screening of several anatomies, modern compact and lightweight robotic arms can further support the physician, i.e. by automatically maneuvering a second imaging probe or tool in order to enable live image fusion [10]. Moreover, such systems eventually incorporate pressure, hand tremor or movement correction [11], or an automatic servoing for tissue biopsy based on a registration to prior US acquisitions [12]. With the goal of fully automatic acquisitions, however, a prior planning of the actual target area of interest is a prerequisite, since whole-body ultrasound scans are impractical and time-consuming. In a clinical interventional setup, planning should be performed by the medical expert based on anatomical data, as given by MRI and CT. Given an appropriate target, (several) 3D ultrasound datasets can then be acquired autonomously during an intervention, without requiring the presence of medical staff guiding the robot. We acknowledge that there have been attempts in the past to incorporate tomographic image information into robotic systems to improve the visualization of ultrasound image information to physicians. However, to our knowledge an integration of these data to enable planning and the *automatic* acquisition of 3D US datasets has not been considered so far.

In this work, we aim at closing this gap in the workflow of current robotic support systems in order to allow for fully automatic 3D ultrasound acquisitions using a robotic imaging system. We present the path towards an autonomous robotic ultrasound system to assist clinicians during interventions by performing multiple and especially reproducible examinations based on pre-interventional planning. To this end, we introduce a first concept for a robotic ultrasound system consisting of a lightweight robotic arm, a clinical ultrasound machine, and a RGB-D camera which is used to observe the scene. This enables the direct planning of a patient-specific trajectory by selecting its start- and endpoints in the patient's MR (or CT) image. A registration of the actual patient position to the MRI allows for the automatic acquisition of 3D ultrasound data. By using intensity-based image registration, we can close the loop and perform an online-update of the patient-to-world registration, accounting for inaccuracies of the RGB-D information as well as patient movement. Thus, the overall goal is to perform fully autonomous image acquisitions within a closed control loop by utilizing 3D surface information, pressure estimations of the robotic system, and image-based servoing, to image regions of interest defined by pre-interventional data[1]. The workflow showing the main steps of the proposed solution is depicted in Fig. 2.

The remainder of this work is organized as follows: Section II will give an overview of related work in the field of autonomous robotic ultrasound acquisitions. We introduce our proposed system and all corresponding and necessary steps in Section III, before experiments and results are presented in Section IV. Finally, opportunities and challenges associated to the system are discussed in Section V, before we draw final conclusions in Section VI.

[1]Supplementary video material to this manuscript features a brief overview of the acquisition setup and workflow with the distinct steps.

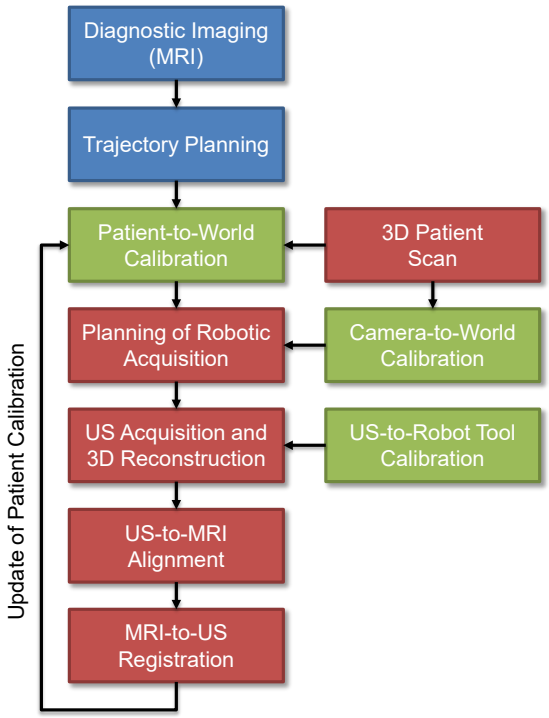

Fig. 2. The proposed workflow for autonomous acquisitions includes a planning based on diagnostic images (blue), calibration of both robot and patient to the world reference (green), and interventional acquisition (red). During acquisition, a 3D patient scan using a structured-light 3D scanner is used for the initial patient to world alignment and robot to world calibration. Following the autonomous robotic ultrasound acquisition, a US to MR registration is conducted to refine the patient to world alignment.

II. Related Work

Over the past two decades, research groups have focused on improving ultrasound image acquisition systems in terms of accuracy, usability, remote controllability as well as by a synchronization with pre-operative planning. For this purpose, (medical) robots have been introduced for different applications. Inspired by [8], we differentiate between situation aware robotic systems, and surgical assistant systems.

Surgical assistant systems primarily serve as advanced tools and are directly controlled by the surgeon, mostly acting in a master-slave fashion. Thereby, these systems are used to augment the operator's ability to perform a specific task, e.g. by allowing precise, jitter-free movements, or the application of constant desired pressure. Systems targeted to specific applications, such as carotid imaging [13], have been explored recently. Furthermore, ultrasound has been added as a tool to the more general surgical da Vinci platform for several applications [14], [15] using specialized drop-in ultrasound probes which can be picked up by the surgeon intra-operatively. Following this concept, a general system providing tele-operated ultrasound using different robotic arms is currently being developed [16]. For a more detailed overview of surgical assistants in ultrasound, the reader is referred to [8].

In contrast to these systems, situation aware robotic system perform at least one task autonomously based on a predefined situation-specific model, requiring an awareness of the task and its surrounding environment. In the following, we will focus our review on this area, as it is the primary target area of this work. In order to do so, we differentiate between Automatic Robotic Support Systems (ARSS) providing automatic support for a defined task, as well as Automatic Data Acquisition Systems (ADAS), targeting fully automatic robotic imaging.

A. Automatic Robotic Support Systems

Boctor *et al.* [17] originally proposed a dual robotic arm setup holding a needle in one arm and a US probe in the second, in order to facilitate an accurate placement of the needle to perform liver ablations. In a later work [18], the robotic arm holding the US probe was replaced by a freehand 3D ultrasound setup to improve general usability while still enabling accurate placement. With respect to an integration of tomographic image data into the robotic ultrasound environment, Zhang *et al.* [19] propose a system combining a robotic arm with ultrasound, optical tracking and MRI data manually registered to the system in order to improve the accuracy of needle placement. Targeting at optimal robotic ultrasound acquisitions to cover a predefined volume of interest, a planning-framework is presented in [20] to perform automatic robotic imaging trajectories. Other groups have focused more on a direct integration of the resulting images into the robot control, facilitating visual servoing based on robotic

ultrasound. Abolmaesumi *et al.* [21] were among the first exploring the combination of robotic arms and ultrasound imaging based on visual servoing by combining an experimental robotic arm holding an ultrasound probe with vessel feature tracking to follow vessels during the ultrasound scan. Krupa *et al.* [11] further explored visual servoing towards an application to ultrasound by using ultrasonic speckle as a main characteristic feature to compensate for potential patient motion both in-plane and out-of-plane with respect to the ultrasound probe. These concepts were later adapted to fully automatic needle insertion under ultrasound guidance [22] and automatic probe positioning in order to allow for an optimal skin coupling with respect to the current position [23]. Recently, a servoing approach using live ultrasound to ultrasound registration was presented, targeting at screw placement in spine surgery [12], [24]. Another intra-operative application of robotic ultrasound was presented for the mobilization and the resection of the internal mammary artery in the course of coronary artery bypass grafting [25], following a servoing approach using color Doppler ultrasound images. For a comprehensive and detailed background of visual servoing, the reader is further referred to [26], [27], where concepts and basic approaches are explained in detail.

B. Automatic Data Acquisition Systems

While all approaches described above perform ultrasound acquisitions semi-automatically or under direct physician guidance, they either require a manual positioning of the US probe, or a proper definition of the working area. A first step towards fully autonomous scanning of ultrasound trajectories [28] combines a designed pneumatic probe holding case with optical tracking in order to enable scanning of a small volume of interest. While this approach has a high potential for certain applications, it still requires manual placement of the autonomous holding cage on the patient and does not enable fully automatic scanning. Using two robotic arms holding separate ultrasound probes, a system for 3D robotic tomographic imaging pioneers the idea of localized 3D tomographic ultrasound imaging [29]. With respect to an intra-operative setting, however, the positioning of two robotic arm seems challenging, which is why the major field of applications for this technique might lie in the imaging of smaller structures. Focusing on vascular applications, a targeted scanning system for the automatic acquisition of vascular 3D ultrasound is presented in [30], combining pressure compensation with vessel tracking mechanisms. Finally, another system focusing on automatic data acquisition for liver screening was recently presented [31], combining RGB cameras for an automatic determination of liver scanning regions, followed by an automated acquisition workflow using different scanning directions. While this system operates fully automatically, it is limited to 2D images only and performs acquisition based on classified body features, which are prone to an erroneous feature determination, limiting the practical applicability of the system.

With this work, we focus explicitly on a system which enables fully autonomous robotic 3D ultrasound acquisitions based on an individual planning from tomographic image data (performed before the intervention). Thereby, the system does not rely on an initial probe positioning, manual registration of image data, or feature tracking in ultrasound or RGB-images. Instead, we use an initial patient-to-robot registration to perform 3D US acquisitions, which are then used in order to refine the overall registration and calibration. To do so, an intensity-based image registration of the pre-aligned 3D US and tomographic image data is used to retrieve a refined and updated patient-to-robot registration. Using the robot's built-in force and torque sensors, this further allows ultrasound acquisitions with optimal skin force, providing repeatable 3D ultrasound volumetric datasets.

III. Methods

This section first introduces the main component of the robotic ultrasound hardware setup in III-A. Next, III-B describes the pre-interventional imaging as well as trajectory planning for autonomous acquisitions. Following our proposed application, all required calibration steps and procedures are defined in III-C, before the interventional workflow for robotic acquisitions is explained in detail in III-D, including robotic control strategies, updates, ultrasound acquisitions and refinement of world-calibrations based on US acquisitions.

A. Hardware Setup

Our system consists of a lightweight robot, the ultrasound device, and a structured light RGB-D 3D scanner. While the ultrasound transducer is directly mounted onto the end-effector of the robotic arm, the 3D scanner is attached to the ceiling and serves as a vision system allowing for the direct calibration and registration of all system parts.

1) Robotic Arm: Based on developments of the German space center (DLR) [32], KUKA introduced a robotic platform targeted at direct human-machine interaction, referred to as 'Intelligent industrial work assistant' - iiwa (KUKA Roboter GmbH, Augsburg, Germany). This system consists of a 7 joint robotic arm with corresponding control units and consequently enables one redundant degree of freedom (6+1 in total). As a result of this design, the robot provides dynamic movement and flexible adaption of trajectories to the working environment. With respect to robotic ultrasound, the incorporated high-accuracy torque sensors in each of the seven joints are evenly important, as a robotic ultrasound platform has to be fully compliant to both patient and staff. Based on significant safety measures for collision detection, the robot subsystem is certified for human-machine-collaboration due to the compliance to standards for functional safety. Thus it is considered to be safe for use in direct interaction with humans. Detailed specifications and design choices can be found in [33].

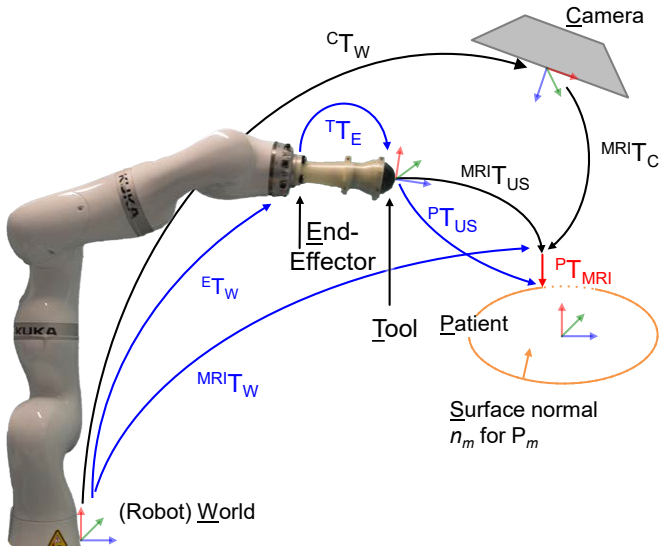

Fig. 3. Using the 3D scan of the patient ${}^{MRI}\mathbf{T}_C$ and the transformation from world to the camera ${}^{C}\mathbf{T}_W$ the world-to-patient transformation ${}^{MRI}\mathbf{T}_W$ is estimated. The tool (transducer apex) reaches the patient's surface (orange) by applying the tool-to-patient transformation under consideration of the surface normals n_m. The discrepancy between the estimated position and the real position of patient is indicated by ${}^{P}\mathbf{T}_{MRI}$, which is detected by intensity-based registration (red).

The KUKA native Sunrise.OS and its Sunrise.Connectivity module allow for the full low-level real-time control of the KUKA iiwa via UDP at rates up to 1 kHz, acting similar to the Fast Research Interface (FRI) [34] as proposed for the previous generation of the KUKA LWR robot arms. In this work, a publicly available software module[2] developed in our lab is utilized to enable a direct interaction between Sunrise.OS and the Robot Operating System[3] (ROS) framework. By doing so, low-level functionality and control provided by the robot manufacturer can be integrated with RGD-D information and high-level robotic interaction through ROS, as required for the proposed system.

2) Structured-Light 3D Scanner: The Kinect 3D camera (Microsoft Corporation, Redmond, Washington, USA) allows for the recording of color images and depth data (RGB-D) at 30 frames per second. The system contains one camera for color video, an infrared laser projector for the projection of a structured light pattern and a camera to record said pattern. This device has been used for several medical applications, such as for a touch-less interface [35], for image-to-patient registration [36], or to provide initialization to US/MRI registration [37]. In the proposed system setup, the camera is attached to the ceiling, focusing on the robotic system and examination bed. The RGB-D data is used in order to i) align the camera within the world coordinate system (camera-to-world calibration - Sec. III-C2), and to ii) register the patient lying on the examination bed to the world reference frame (patient-to-world calibration - Sec. III-C3).

3) Ultrasound System: In general, any ultrasound device could be incorporated within an autonomous system by using the video output of an arbitrary clinical US device and recording those images using hardware frame-grabbers connected to the device. In practice, such systems provide a partially diminished image quality due to compression artifacts and - more importantly for robotic acquisitions - a temporal lag between the acquired image by the US machine and the recorded frame via frame-grabber. To enable direct acquisitions and advanced servoing in real-time, we favor ultrasound devices providing direct interfaces for retrieving acquired ultrasound data with minimal temporal delay. An Ultrasonix RP ultrasound machine (Analogic Corporation, Peabody, Massachusetts, USA) is used in combination with a 4DC7-3/40 4D Convex curvilinear transducer used in single-plane mode. The system provides the ulterius[4] API, enabling both the streaming of ultrasound data to a client as well as the control of US acquisitions parameters remotely through the Ethernet connection.

B. Diagnostic Imaging and Trajectory Planning

The goal of an autonomous 3D ultrasound acquisition is to perform a robotic-ultrasound trajectory based on a planning on pre-interventional images. With respect to a fast and intuitive planning, a physician wants to see either i) the region around a specific target, such as for liver or lymph node biopsy, or ii) scan a defined region, such as an organ, in order to perform a diagnosis based on these images. In this work we focus on the general case, where the operator can select the region of

[2]https://github.com/SalvoVirga/iiwa_stack
[3]http://www.ros.org/
[4]http://www.ultrasonix.com/wikisonix/index.php/Ulterius

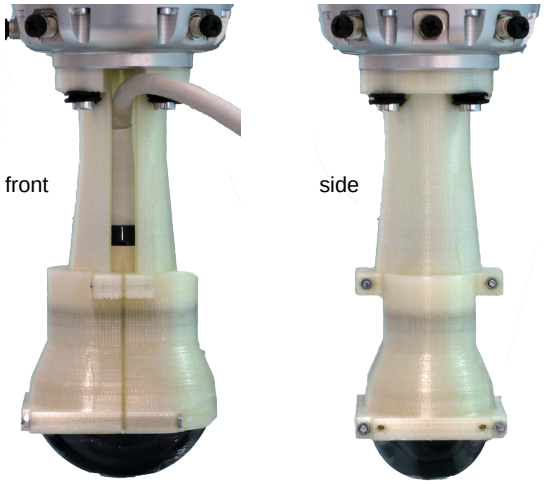

Fig. 4. Ultrasound probe mount designed for lightweight arm end-effector.

interest (ROI) directly. For our setup, a T2-weighted MRI volume is used as basis to determine the ROI for the ultrasound scan. The physician simply selects the start- and endpoint P_s, P_e of a trajectory in the MRI data, where the acquisition path direction $\vec{d_t}$ is defined by these points:

$$\vec{d_t} = P_e - P_s. \tag{1}$$

To enable both the transfer of the planned acquisition path to the robot world coordinate system and the initialization of the MRI/US registration, the patient's surface is extracted from the MRI by thresholding the image. Based on the patient-to-world calibration (see Section III-C3), the trajectory direction and points can be directly transformed into world coordinates. Consequently, the segmented patient surface is exploited to determine the points of interest for the autonomous ultrasound scan.

C. System Calibration

1) US Tool Configuration and Calibration: The transformation ${}^{\mathrm{T}}\mathbf{T}_{\mathrm{E}}$ from the robot end-effector to the transducer tool tip can be obtained directly from the CAD models used for manufacturing of a customized mount for the rigid attachment of the US probe to the robot end-effector. The probe mount utilized for the robotic system is shown in Fig. 4 and is tailored to the ultrasound transducer for optimal and rigid attachment. To perform ultrasound acquisitions, a second spatial transformation ${}^{\mathrm{US}}\mathbf{T}_{\mathrm{T}}$, pointing from the probe tool tip to the US image origin, is defined by the image scaling (transformation from px to mm) and size with respect to the US probe apex

$${}^{\mathrm{T}}T_{\mathrm{US}} = \begin{bmatrix} s_x & 0 & 0 & s_x t_x \\ 0 & s_y & 0 & s_y t_y \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \tag{2}$$

where s_x, s_y determine the resolution of the ultrasound image, and t_x, t_y the translation from the apex, defined by the center of curvature, to the image origin. If an additional refinement of the transformation from the ultrasound image to the transducer apex is required, a standard ultrasound to tracking calibration technique can be applied [38].

As the used lightweight robot provides force sensors in each joint, an additional calibration with respect to the weight and center of mass of the tool is necessary to allow the force controlled motion. For an accurate load calibration, a proprietary calibration algorithm provided with the robot is used to determine these values in 3D.

2) Camera-to-World calibration: This calibration step allows for the control of the robot - representing the world reference frame - as observed in camera coordinates. In the case of RGB-D cameras, this transformation relates from the camera to the robot arm, and can be computed either by using the 2D RGB images or the 3D information obtained by the depth sensor. An accurate calibration can be achieved with both techniques, although the 3D case requires more effort, user interaction and processing time [39]. Alternately, in [20], an additional optical tracking system is used to achieve a good calibration between robot and camera. However, by leveraging the 2D RGB images, it is possible to compute the transformation ${}^{\mathrm{C}}\mathbf{T}_{\mathrm{W}}$ between the robot base and the RGB-D camera by using a standard hand-eye calibration method [40] in its 'eye-on-base' variant, without requiring an additional optical tracking system. To do so, an Augmented Reality (AR) marker is placed on the robot flange and the transformation between the marker and the camera estimated using the ArUco [41] tracking library. Several tracked

poses combining random (non co-linear) movements around several rotation axis are recorded and used as correspondence pairs for calibration. The robot forward kinematics and the presented transformation ${}^{\mathrm{T}}\mathbf{T}_{\mathrm{E}}$ complete the chain of transformations from the camera reference frame to the US transducer. It should be noted that by using this calibration, the exchange of the US transducer with another one does not require a full camera recalibration, but only recomputing the last transformation of the chain between the end-effector and the US probe ${}^{\mathrm{T}}\mathbf{T}_{\mathrm{E}}$. Furthermore, the AR marker-based calibration is used for the camera-to-world calibration only, allowing for a fully marker-free data acquisiton and update. The transformation between the RGB and the depth sensors of the Kinect can be obtained as proposed by [42]. For the hardware used, this relates to a translation of 2.5 cm along the transverse axis of the camera.

3) Patient-to-World Calibration: With respect to a fully autonomous acquisition, it is necessary to transfer trajectories planned in a tomographic image to the actual position and orientation of a patient lying on the bed. To do so, acquired 3D RGB-D information of the patient (Sec. III-A2) provides surface information, which can be extracted simultaneously from the tomographic images. These data can then be used directly for aligning both surfaces and consequently determining their spatial transformation. In the following section, we will describe the necessary steps to enable such a calibration in detail.

Surface Extraction from MRI Data: To extract the patient (or phantom) skin surface from the MRI images, we employ an approach which does not require complex segmentations, as the extracted and matched surfaces will be refined throughout the acquisition process using image-based registration. Therefore, tomographic data (in this work T2-weighted MRI) is thresholded first, neglecting all values $I < \tau$. Throughout all our experiments, $\tau = 100$ provided a sufficient discrimination between the patient and the background (acquisition table and room around patient). As the resulting mask contains holes and is partially irregular, a morphological closing is performed on the thresholded image containing of subsequent dilation and erosion. The surface can then be retrieved from the image data [43], where only the connected surface component covering the highest ratio of total pixels is kept for further processing

$$\Omega_B = \underset{i}{arg\,max} \sum \Omega_i, \quad \Omega_i = \{x_1, x_2, ..., x_m\}, \tag{3}$$

with x_m being the surface positions of the points contained in the component Ω_i in 3D space.

Spatial Change Detection in RGB-D Data: To separate the patient from the background and other objects within the view of the camera, a spatial change detection is performed. Octrees are often used in computer graphics to partition a volume by recursively dividing cubes [44]. First, a background point cloud is recorded and added to an octree. After positioning the patient on the bed, a second point cloud is retrieved and added to another octree. The differences in the tree structure allow the efficient determination of newly added points, which represent the background-subtracted object, in our case the patient. To increase robustness and compensate for noise, a minimum of n pixels (experimentally set to 2) must be contained in each tree node.

Surface Matching: The alignment of surfaces can either be achieved using Iterative Closest Points (ICP), a non-linear derivative of ICP, or by directly aligning features. As the patient surface shape deviates strongly from the shape of the background (e.g. table or ultrasound scanner) a feature alignment process is applicable. On the foundation of the calibration refinement using intensity-based registration (Sec. III-D5), our framework automatically accounts for local deformations and inaccuracies. Consequently, we employ ICP as surface matching method, as it provides a robust and especially highly efficient global alignment, which is then used to initialize the intensity-based registration.

Result of the surface-matching between extracted MRI and RGB-D information will be a rigid transformation from MRI-space to RGB-D camera space

$$P_i^C = {}^C\mathbf{T}_{MRI} P_i^{MRI} = ({}^{MRI}\mathbf{T}_C)^{-1} P_i^{MRI}, \tag{4}$$

with P_i^{MRI}, P_i^C being the surface points in MRI and camera-space respectively.

D. Autonomous US Acquisition

Following the workflow as shown in Fig. 2, this section will describe all steps carried out to perform one or multiple autonomous acquisitions.

1) Planning of Robotic Acquisition: Based on the alignment of both camera and patient within the world reference (Sec. III-C), the next step is to transfer the previously planned image trajectory to a robotic control trajectory which can be executed by the robotic arm. To allow accurate and safe trajectories, a proper force control strategy of the robotic manipulator, and the planning of the acquisition are necessary.

Stiffness and Force Control: The manipulator is mainly operated using a stiffness controller, which represents a Cartesian virtual spring between the desired (setpoint) position x_{set} and current (measured) position x_{msr}. The control law for this strategy is defined by the transposed Jacobian matrix J^T

$$\tau_{Cmd} = J^T(k_c(x_{set} - x_{msr}) + D(d_c)) + f_{dyn}(q, \dot{q}, \ddot{q}), \tag{5}$$

where k_c is the Cartesian stiffness of the virtual spring $k_c(x_{set} - x_{msr})$. The damping term $D(d_c)$ is dependent on the normalized damping value, while the dynamic model of the manipulator $f_{dyn}(q, \dot{q}, \ddot{q})$ is predefined by the manufacturer [34]. The resulting joint torque is computed by a Cartesian law using J^T.

In order to allow for compliant force applied in the US probe direction, the force control is modified such that the force term $D(d_c)$ is set to a high damping in the respective probe direction, allowing for a compliant movement. Thus the high accuracy torque sensors and real-time forward kinematics computation of the system are exploited in this view in order to provide an acoustic force coupling without applying excessive forces (> 25 N) on the skin. The stiffness controller allows to achieve an indirect force control [45], considering the relationship between the deviation of the end-effector position and orientation from the planned motion as well as the contact force and moment. The online parametrization of the virtual spring values makes it possible to maintain the contact force constant. The Cartesian stiffness along the direction of the US probe is set in the range [125-500] N/m according to the anatomy of the patient, and the force desired is parametrized as 5 N. The stiffness and forces in the other directions are parametrized to 2000 N/m and 0 N accordingly. If an excessive force > 25 N occurs, the robot's internal subroutine will automatically stop the acquisition process. This feedback loop enables the compliant motion constrained by the patient or other objects.

Planning of Acquisition Path: Based on the selected points (see Section III-B) and corresponding surface normals provided by the RGB-D data, the acquisition path is planned. The selected start and end points of the trajectory P_s,P_e define the direction vector of the trajectory (see Eq. (1)), which is used to define equidistant sampling points $Q^1, Q^2, \ldots, Q^i$ along the line with a distance of 2 cm. For each sampling point, the closest surface point $Q^i \rightarrow P^k$ is retrieved by a nearest-neighbor search. Along with the corresponding surface normal directions n_k, the points can then be used directly as target tool poses for robotic movements. It should be noted that the force controller implicitly commands the robot to adapt the tool position until the surface is reached. As the robot approaches the next trajectory point P^{k+1}, the direction of the transducer is stepwise changed to the subsequent surface normal position defined by P^k, P^{k+1} by the robot control.

2) US Acquisition and 3D Reconstruction: The live streams of robot tracking data and ultrasound images are transmitted via Ethernet, therefore the offset between tracking and US data is small compared to the framework overhead. In order to enable 3D registration with the diagnostic image, a volume compounding is performed using a backward normalized-convolution approach following [46], yielding regular-spaced 3D volume data from the arbitrarily sampled ultrasound frames in 3D space.

3) US-to-MRI Alignment: Making use of all previously estimated transformations from camera to MRI ${}^{MRI}\mathbf{T}_C$, world to camera ${}^{C}\mathbf{T}_W$ and ultrasound to tool ${}^{T}\mathbf{T}_{US}$, we can transform both ultrasound and MRI data into the world space using the respective transformations from both image spaces to the world reference frame

$$ {}^{W}\mathbf{T}_{MRI} = ({}^{MRI}\mathbf{T}_W)^{-1} = ({}^{MRI}\mathbf{T}_C \cdot {}^{C}\mathbf{T}_W)^{-1} \tag{6} $$

$$ {}^{W}\mathbf{T}_{US} = {}^{W}\mathbf{T}_E \cdot {}^{E}\mathbf{T}_T \cdot {}^{T}\mathbf{T}_{US}. \tag{7} $$

Both image datasets will be roughly aligned after transformation into the world coordinate frame, relying on prior ultrasound and patient to world calibrations.

4) MRI-to-US Registration: Supposing a rough overlap of MRI and US images as described above, an intensity-based MRI/US registration can directly be initialized in order to obtain an updated transformation between ultrasound and tomographic images.

The LC^2 similarity method allows for the registration of MRI and US images by correlating the MRI intensities and gradient magnitudes to the US intensity values. High robustness, wide convergence range and the application for rigid, affine and deformable registration have been shown in [47]. This approach is utilized in a two-step process. Based on the calibration and transformation chain, we can bring recorded ultrasound volume data directly into the MRI reference frame (or both into the world reference frame)

$$ {}^{MRI}\mathbf{T}_{US} = {}^{MRI}\mathbf{T}_W \cdot {}^{W}\mathbf{T}_{US} = ({}^{W}\mathbf{T}_{MRI})^{-1} \cdot {}^{W}\mathbf{T}_{US}. \tag{8} $$

In a second step, LC^2 is used to determine the rigid transformation component, aligning the transformed US and MRI images in world space in reference to the actual patient position

$$ {}^{P}\mathbf{T}_W = {}^{P}\mathbf{T}_{MRI}\,{}^{MRI}\mathbf{T}_W, \tag{9} $$

with ${}^{P}\mathbf{T}_{MRI}$ being the updated transformation to compensate for patient movement, deformation, as well as tracking and detection inaccuracies of ${}^{C}\mathbf{T}_{MRI}$ with respect to the world reference frame, *c.f.* Eq. (7). Optionally, a second step consisting of an affine and deformable registration using the same similarity measure can be performed, resulting in a precise and direct correspondence between voxels in the compounded 3D-US and reconstructed MRI volumes, which finally closes the loop of the autonomous ultrasound acquisition system and enabling MRI guided interventions. We evaluate both rigid and affine registrations; the decision whether both are required relies on the specific application in mind. For instance, if bony structures are to be scanned rather than soft tissue, a rigid alignment would be sufficient based on our experience.

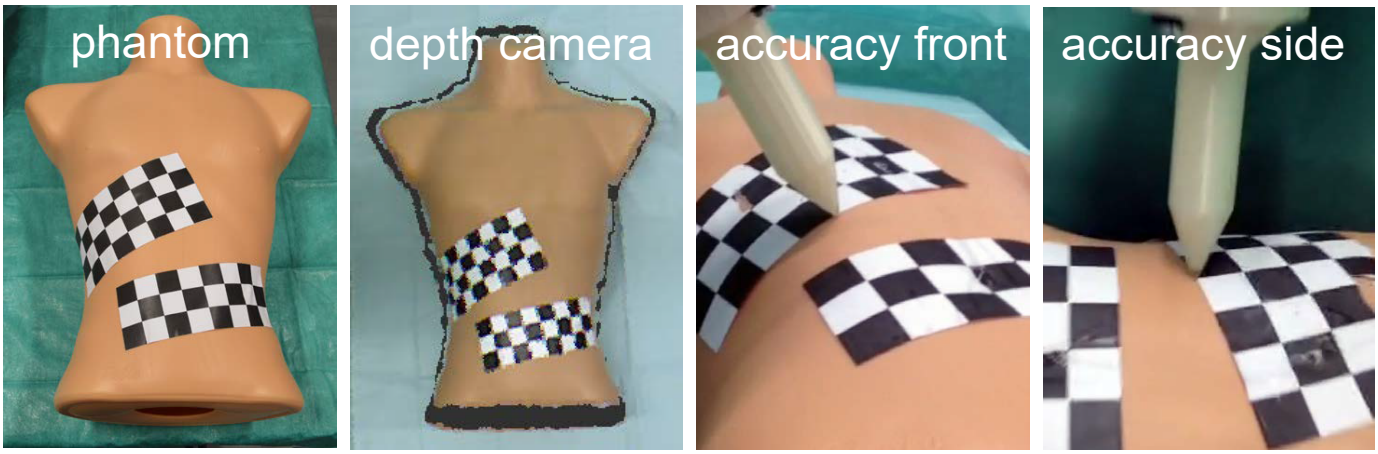

Fig. 5. From the colored point cloud (depth camera), individual checkerboard corner points on the uneven surface are manually selected as target points. The points can be targeted with an accuracy of 2.46 mm in x-y-direction, and 6.43 mm in z-direction.

5) Update of Patient Calibration: Based on the estimated transformations from the (robotically) acquired 3D ultrasound data and the MRI aligned to the world, the transformation from tomographic imaging to ultrasound space can be refined, in order to enable a more precise initialization of subsequent acquisitions. This becomes especially interesting as such acquisitions can be performed for an on-line refinement of the whole system calibration through an image-based feedback and update. In an interventional setup, an initial acquisition is performed at the beginning to optimize the robot- and patient-to-world calibrations. Subsequent planned trajectories can then be performed automatically based on an updated calibration. Making use of the estimated rigid transformation ${}^{P}\mathbf{T}_{MRI}$ aligning the US and MR volumes in world space as described above, the patient-to-world calibration can be updated accordingly to

$$ {}^{P}\mathbf{T}_{US} = {}^{P}\mathbf{T}_{MRI} \cdot {}^{MRI}\mathbf{T}_{C} \cdot {}^{C}\mathbf{T}_{W} \cdot {}^{W}\mathbf{T}_{US}. \tag{10} $$

On the one hand, this reduces processing time, as no rigid alignment is required for those US volumes. On the other hand, the comparability of subsequent records is fully maintained by this approach, as images are provided for the identical anatomy and planned trajectory. For an exemplary application of US-guided liver biopsy, several acquisitions of the liver could be conducted throughout the procedure, enabling a reliable 3D image-guidance during needle insertion.

IV. Experiments and Results

A. Robot/Camera Calibration and Robot Control

Utilizing the depth images augmented with color information, we evaluate the accuracy of the tool and camera-to-world calibrations by equipping the robot manipulator with a rigid acuminate tool and moving the tool tip onto selected points on a realistic rigid surface. In this view, after tool configuration, the camera-to-robot calibration is performed to obtain the required transformations as described in Sec. III-C2, where 13 poses are used in total for the calibration. Next, the RGB-D data of an upper torso phantom surface is recorded and the surface normals computed. The planned acquisition path is defined by manually selecting multiple intersection points on a printed checkerboard attached to the phantom's curved surface, see Fig. 5. The robot is commanded in position control mode to move onto the individual points and the distances to the actual intersection points are measured manually with a caliper. The experiment is performed twice by selecting 10 points in each run and recalibrating the whole system in between the two sessions. For the first run, the results yielded an average accuracy $|x_{set} - x_{msr}|$ of 2.42 ± 1.00 mm on the x-y plane and 7.20 ± 3.30 mm along the z-axis (camera's depth axis). For the second one, the accuracy was estimated as 2.5 ± 0.90 mm (x-y plane) and 5.64 ± 4.04 mm (z-axis). Overall, accuracies were 2.47 ± 0.96 mm (x-y) and 6.43 ± 3.68 mm (z). During all experiments, the camera was placed at around 1.5 m distance to the phantom. This shows that the calibration yields reproducible results, while the inaccuracies are dependent on the RGB-D camera information. It should be noted that these values are in line with the reported average spatial resolution of the camera, being 3.0 mm (x-y plane) and 10 mm (depth axis) for a camera at 2.0 m from its target [48]. It is also important to notice that while the $x-y$ accuracy directly affects the robot poses, the z errors are effectively compensated by the desired-force control of the robot. In this view, the resulting system layout can compensate for the inaccuracies of the RGB-D camera. Based on these accuracies, the calibration and overall system accuracy is anticipated to be sufficient for the initialization of an image-based registration, and thus also for the full system.

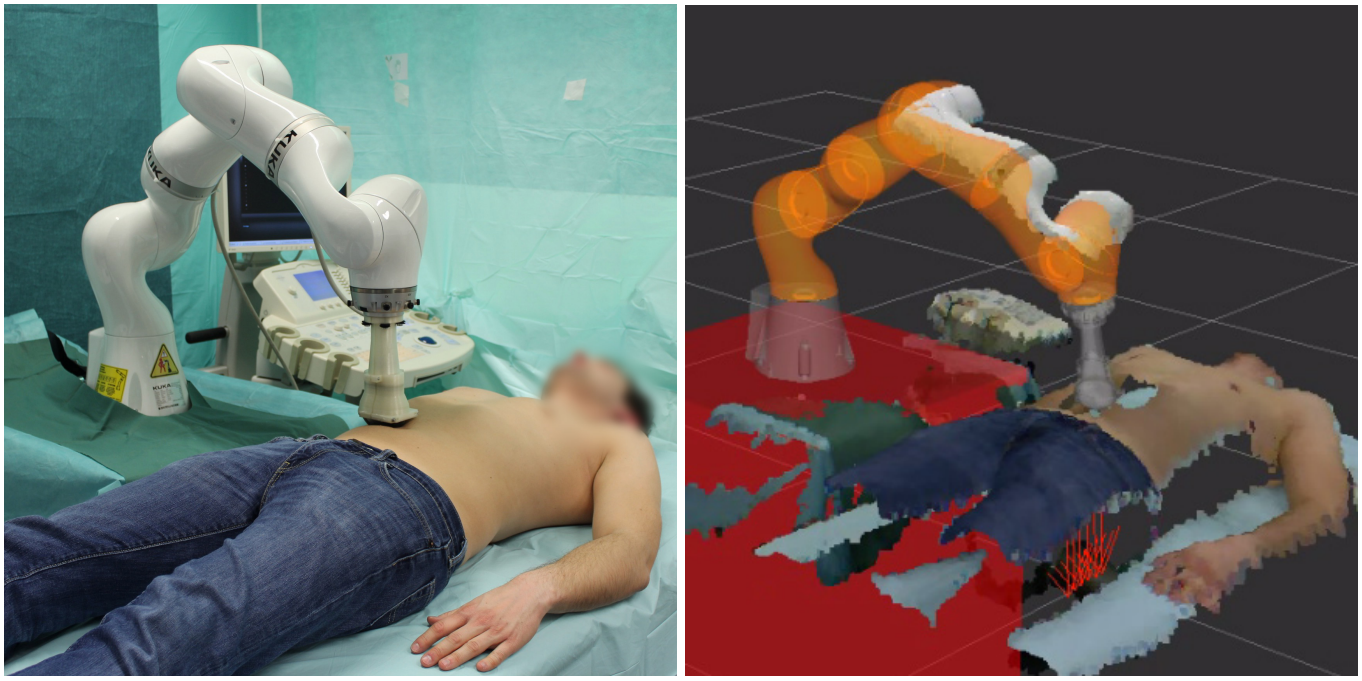

Fig. 6. Scanning setup showing the real system setup scanning a volunteer (left) as well as the visualization of all involved components accordingly with the point cloud given by the RGB-D camera.

B. MRI/US Image Acquisition and Registration

To allow for an evaluation of the overall system, we first use a triple-modality 3D abdominal phantom (Model 057A, Cirs Inc., Norfolk, Virginia, USA), which provides anatomically correct objects such as the liver, cysts, ribs and vertebrae. The phantom is targeted at MRI, CT, and ultrasound imaging and provides realistic image properties for all modalities. We then perform similar automatic robotic US acquisitions on two healthy volunteers, for whom an MRI scan was performed prior to this work.

Following the workflow presented, the MRI images are transformed into the world reference frame by matching the 3D scanned surface with the MRI surface (see Sec. III-C3). Based on this global registration, trajectories are planned in the MRI by selecting start- and endpoint. Figure 6 shows the actual scanning setup for a healthy volunteer, where the robot arm performs an automatic acquisition based on a planned trajectory. Given a perfect system accuracy, i.e. a system with perfect calibration and without any tracking or imaging inaccuracies, an automatic acquisition of a 3D ultrasound volume by the robotic system would yield a perfect overlap between reconstructed 3D US data and tomographic imaging. In reality, however, the image alignment will be affected, which is why an intensity-based registration is then performed in order to detect the offset and correct for system inaccuracies. The alignment by intensity-based registration is first performed first rigidly (accounting for tracking, calibration, patient registration, and movement), followed by an affine improvement (accounting for image scaling as well as patient and tissue movement/deformation). Due to the fixed direction of the trajectory and a constant skin-pressure, an affine registration is sufficient for this application. The resulting rigid transformation directly indicates the accuracy of the overall 3D scanning system based on the alignment of planned and real acquisition (distance between desired and actual US). In a second step, we then use the rigid transformation part of the registration to refine the overall system and robot calibration according to Eq. (9). The experiment is repeated with a second planned trajectory, followed by an evaluation of rigid and affine registrations between the calibration-updated scans and the tomographic data. The hypothesis is that by using the rigid part of the intensity-based registration, we can align the pre-interventional MRI with the interventional US data. This closes the loop between world calibration and robotic-system, where the system calibration can be refined online. In order to also evaluate the errors of the system in different situations, we repeat the overall experiment two times for the phantom, and between each experiment the phantom is moved on the examination bed. For the scans on volunteers, we repeat the overall experiment three times per person, where the person stands up and lies down between each run, evaluating also potential variabilities based on the patient-to-world calibration. Robotic acquisitions are taken at a similar speed as freehand acuqisitions, such that a normal acquisition requires < 30 s. Ultrasound scans are planned in the the abdominal region and contain organs (e.g. liver, kidney) and vessels. Acquisitions may also contain ribs, where the robot follows the natural curvature of the body surface without exceeding the maximal force applied. Ultrasound parameters are set manually, although modern US systems provide automatic adaption allowing for a direct optimization of settings. Coupling gel is also applied manually, and the volunteers must hold their breath during the scan. Safety of acquisitions is guaranteed by the collision detection provided by internal subsystems

incorporated in the KUKA lightweight system, which would stop actions automatically if the force desired is exceeded. The patient-to-world calibration is performed once before each acquisition, but could be also updated online in future. The global alignment between the patient and MRI surface requires less than 10s, and the intensity-based refinement of the registration requires less than 30s to complete.

For all acquisitions of the phantom and humans combined, the average translation between the initial US-volume and MRI is 4.47 ± 2.15 mm for iteration one (without calibration-update) compared to 0.97 ± 0.76 mm in the second run after the overall patient calibration is updated. Similarly, the rotational part clearly improves after the update of the calibration ($4.50 \pm 2.24^\circ$ before update, versus $1.08 \pm 0.76^\circ$ after), where the rotational error is determined as the Euclidean norm of the three rotation angles. For an additional affine registration after the initial rigid registration, the average translation results in 1.02 ± 0.77 mm, which shows that the initial registration successfully accounts for potential tracking and point cloud inaccuracies (rotation errors are $1.43 \pm 1.42^\circ$). Resulting ultrasound datasets and registrations are depicted in Fig. 7 for a volunteer acquisition, and all results of the different experiments are listed in Table 1.

	Translation [mm]	
	Phantom	Human
Rigid scan #1	4.69 ± 1.35	4.39 ± 2.47
Rigid scan #2	1.18 ± 0.52	0.89 ± 0.84
Affine scan #1	1.30 ± 1.06	0.92 ± 0.75
	Rotation [$^\circ$]	
Rigid scan #1	3.66 ± 1.99	4.79 ± 2.42
Rigid scan #2	0.40 ± 0.32	1.30 ± 0.74
Affine scan #1	0.46 ± 0.19	1.37 ± 1.60

Table 1. Quantitative results based on registration

Thereby, initial, rigid and affine refer to the initial US scan, the registration based on the calibration update, and the affine registration after rigid registration, respectively.

V. Discussion

Based on the millimeter accuracy achieved on repeated acquisitions with an updated calibration, our results suggest the feasibility of the overall approach as well as a potential path for a clinical integration of the presented system. The initial system-calibration using RGB-D information showed a maximum error below 1 cm, which ensures that the acquired 3D-US datasets will lie within the capture range of the intensity-based registration [47]. Therefore, the system should be able to deal with challenging clinical settings, where a higher deviation might lead to local minima during the registration optimization. As the position and orientation of the acquired ultrasound images are aligned with the diagnostic image, US data can be correlated directly to the underlying tomographic image data, aiding the medical expert in terms of ultrasound visualization. The results of the intensity-based registration (Table 1) show the overall translation/rotation of the US volume with respect to the MRI data, where the affine registration of the second scan appeared to perform worse than the rigid registration after calibration-refinement. While the quantity of data is too little to allow for a statistical meaningful judgement, the local deformation of tissue by the robotically held US probe is the primary reason for this discrepancy in our experience. In view of a specific clinical application in mind, it is thus important to analyze and adapt the utilized registration tools further, especially with respect to a deformable registration of soft tissue.

The integration of ROS and Sunrise.OS as used in this work allows for a full control of the KUKA APIs directly from components being available within ROS. Furthermore, a direct access to the robot state is provided by the KUKA APIs, including native features such as self collision avoidance. In this regard, the utilized approach allows for the best combination of both worlds. This facilitates the rapid development of new approaches, but also ensures safety of the robot environment, i.e. by collision detection and emergency halt features.

In view of a clinical integration, it should be noted that our results provide only a first feasibility evaluation, where future work clearly needs to focus not only on healthy volunteers but also diseased and pathological anatomy. Methods such as patient surface extraction and dynamic registration will require further adaptions, aiming at a clinical integration in the future. Our experiments also showed that the currently used examination bed is not optimal for the robot working space, as the end-effector is moving almost at the base of the robot. With respect to a clinical application, we thus suggest a height-adjustable bed, which could be directly integrated with the imaging system, such that the optimal height would be adjusted based on the planned acquisition. Besides that, acoustic coupling between the US probe and patient surface needs to be explored, facilitating US scans with either automatic gel application, or the exploration of other alternatives.

To this end, our future work will also focus on the possible online-optimization of the trajectory based on the initial planning, as we partially experienced suboptimal US image quality for the selected trajectories in this work. In this view, also an automatic change of the applied coupling force with respect to the target surface (e.g. fat vs. muscle) is analyzed in ongoing efforts.

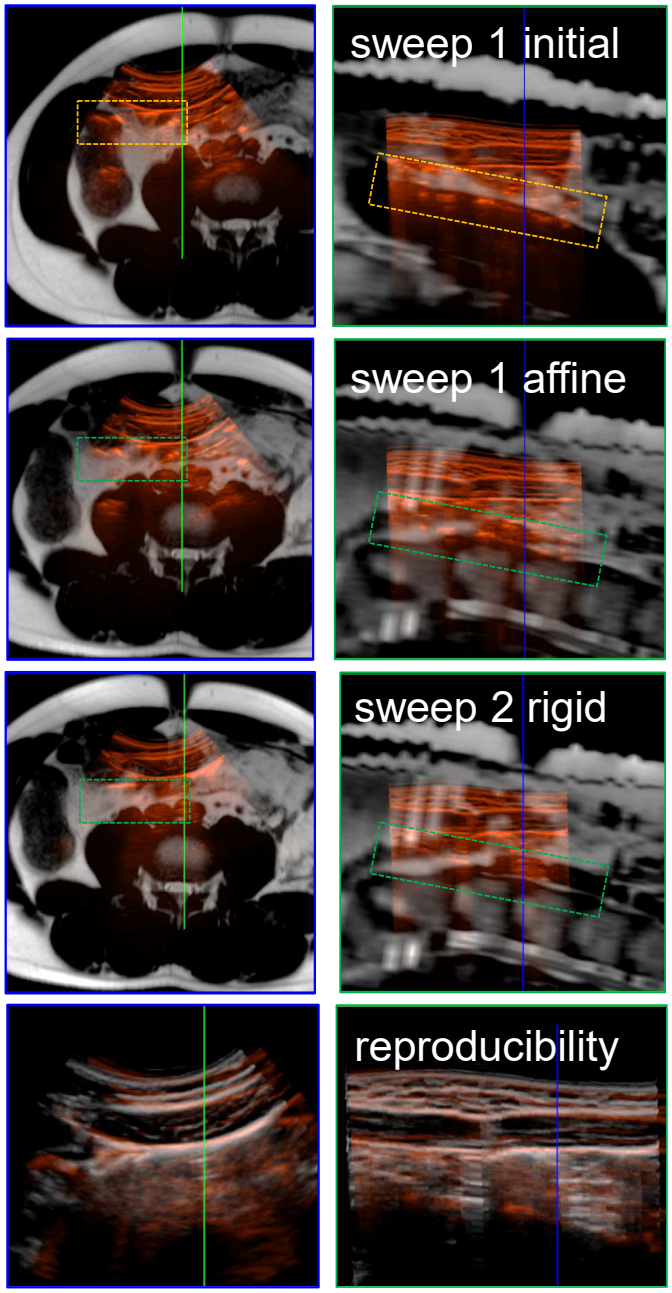

Fig. 7. Left column: axial plane; right column: saggital plane. The initial calibration shows partial misalignment (orange box), while the registrations for sweep #1 and #2 account for this offset. An overlay of both scans also shows the high reproducibility of the sweeps.

VI. Conclusion

We have presented a path to an autonomous robotic ultrasound system in order to enable imaging and support during interventions. The set of methods presented shows the basic feasibility of the automatic scanning framework, allowing fast and efficient robotic 3D ultrasound acquisitions based on pre-interventional planned trajectories. On the foundation of an integrated system consisting of a baseline tomographic image with 3D RGB-D information, we automatically register patient data and perform automatic robotic acquisition. We introduced a closed-loop calibration update based on image-based registration to facilitate the acquisition of reproducible 3D US data. Our results show that the overall accuracy of the system is sufficient for clinical applications. Despite challenges which need to be overcome before such systems could be used in daily routine, this work will hopefully facilitate the clinical acceptance of automatic and autonomous robotic ultrasound scanning systems in the future.

ACKNOWLEDGMENT

We thank the department of nuclear medicine at Klinikum rechts der Isar for the MRI acquisitions, as well as ImFusion GmbH for the provided registration software. This work was partially supported by the Bayerische Forschungsstiftung project RoBildOR.

REFERENCES

[1] L. Beales, S. Wolstenhulme, J. Evans, R. West, and D. Scott, "Reproducibility of ultrasound measurement of the abdominal aorta," *British Journal of Surgery*, vol. 98, no. 11, pp. 1517–1525, 2011.

[2] A. Gee, R. Prager, G. Treece, and L. Berman, "Engineering a freehand 3D ultrasound system," *Pattern Recognition Letters*, vol. 24, no. 4, pp. 757–777, 2003.

[3] O. V. Solberg, F. Lindseth, H. Torp, R. E. Blake, and T. A. Nagelhus Hernes, "Freehand 3D ultrasound reconstruction algorithmsa review," *Ultrasound in medicine & biology*, vol. 33, no. 7, pp. 991–1009, 2007.

[4] M. J. Monaghan, "Role of real time 3D echocardiography in evaluating the left ventricle," *Heart*, vol. 92, no. 1, pp. 131–136, 2006.

[5] H. N. Cardinal, J. D. Gill, and A. Fenster, "Analysis of geometrical distortion and statistical variance in length, area, and volume in a linearly scanned 3-D ultrasound image," *Medical Imaging, IEEE Transactions on*, vol. 19, no. 6, pp. 632–651, 2000.

[6] P. De Rango, "Aneurysm diameter measurement: A challenging and frustrating task," *European Journal of Vascular and Endovascular Surgery*, vol. 43, no. 1, p. 34, 2012.

[7] A. Fenster, D. B. Downey, and H. N. Cardinal, "Three-dimensional ultrasound imaging," *Physics in medicine and biology*, vol. 46, no. 5, p. R67, 2001.

[8] A. Priester, S. Natarajan, M. Culjat *et al.*, "Robotic ultrasound systems in medicine." *IEEE Trans Ultrason Ferroelectr Freq Control*, vol. 60, no. 3, pp. 507–523, 2013.

[9] P. Chatelain, A. Krupa, and N. Navab, "3D ultrasound-guided robotic steering of a flexible needle via visual servoing," in *Robotics and Automation (ICRA), 2015 IEEE International Conference on*. IEEE, 2015, pp. 2250–2255.

[10] M. Esposito, B. Busam, C. Hennersperger, J. Rackerseder, A. Lu, N. Navab, and B. Frisch, "Cooperative robotic gamma imaging: Enhancing us-guided needle biopsy," in *Medical Image Computing and Computer-Assisted Intervention*. Springer, 2015, pp. 611–618.

[11] A. Krupa, G. Fichtinger, and G. D. Hager, "Real-time motion stabilization with b-mode ultrasound using image speckle information and visual servoing," *The International Journal of Robotics Research*, vol. 28, no. 10, pp. 1334–1354, 2009.

[12] R. Kojcev, B. Fuerst, O. Zettinig, Fotouhi, S. Lee, R. Taylor, E. Sinibaldi, and N. Navab, "Dual-robot ultrasound-guided needle placement: Closing the planning-imaging-action loop," *Information Processing in Computer Assisted Interventions - IPCAI*, June 2016.

[13] R. Nakadate, Y. Matsunaga, J. Solis, A. Takanishi, E. Minagawa, M. Sugawara, and K. Niki, "Development of a robot assisted carotid blood flow measurement system," *Mechanism and machine theory*, vol. 46, no. 8, pp. 1066–1083, 2011.

[14] S. Billings, N. Deshmukh, H. J. Kang, R. Taylor, and E. M. Boctor, "System for robot-assisted real-time laparoscopic ultrasound elastography," in *SPIE Medical Imaging*. International Society for Optics and Photonics, 2012, pp. 83 161W–83 161W.

[15] C. Schneider, J. Guerrero, C. Nguan, R. Rohling, and S. Salcudean, "Intra-operative Pick-up ultrasound for robot assisted surgery with vessel extraction and registration: A feasibility study," in *Information Processing in Computer Assisted Interventions - IPCAI*. Springer, 2011, pp. 122–132.

[16] F. Conti, J. Park, and O. Khatib, "Interface design and control strategies for a robot assisted ultrasonic examination system," in *Experimental Robotics*. Springer, 2014, pp. 97–113.

[17] E. M. Boctor, G. Fischer, M. A. Choti, G. Fichtinger, and R. H. Taylor, "A dual-armed robotic system for intraoperative ultrasound guided hepatic ablative therapy: a prospective study," in *Robotics and Automation, 2004. Proceedings. ICRA'04. 2004 IEEE International Conference on*, vol. 3. IEEE, 2004, pp. 2517–2522.

[18] E. M. Boctor, M. A. Choti, E. C. Burdette, and R. J. Webster Iii, "Three-dimensional ultrasound-guided robotic needle placement: an experimental evaluation," *Int J Med Robot Comp*, vol. 4, no. 2, pp. 180–191, 2008.

[19] D. Zhang, Z. Li, K. Chen, J. Xiong, X. Zhang, L. Wang *et al.*, "An optical tracker based robot registration and servoing method for ultrasound guided percutaneous renal access," *Biomedical engineering online*, vol. 12, no. 1, p. 47, 2013.

[20] C. Graumann, B. Fuerst, C. Hennersperger, F. Bork, and N. Navab, "Robotic ultrasound trajectory planning for volume of interest coverage," *International Conference on Robotics and Automation, IEEE*, May 2016.

[21] P. Abolmaesumi, S. Salcudean, and W. Zhu, "Visual servoing for robot-assisted diagnostic ultrasound," in *Engineering in Medicine and Biology Society, 2000. Proceedings of the 22nd Annual International Conference of the IEEE*, vol. 4. IEEE, 2000, pp. 2532–2535.

[22] A. Krupa, "3D steering of a flexible needle by visual servoing," in *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2014*. Springer, 2014, pp. 480–487.

[23] P. Chatelain, A. Krupa, and N. Navab, "Optimization of ultrasound image quality via visual servoing," in *Robotics and Automation (ICRA), 2015 IEEE International Conference on*. IEEE, 2015, pp. 5997–6002.

[24] O. Zettinig, B. Fuerst, R. Kojcev, M. Esposito, M. Salehi, W. Wein, J. Rackerseder, B. Frisch, and N. Navab, "Toward real-time 3D ultrasound registration-based visual servoing for interventional navigation," *International Conference on Robotics and Automation, IEEE*, May 2016.

[25] F. A. Fröhlich, G. Passig, A. Vazquez, and G. Hirzinger, "Robot assisted internal mammary artery detection for coronary revascularisation surgery," in *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*. IEEE, 2010, pp. 1849–1855.

[26] F. Chaumette and S. Hutchinson, "Visual servo control. I. Basic approaches," *Robotics & Automation Magazine, IEEE*, vol. 13, no. 4, pp. 82–90, 2006.

[27] F. Chaumette, S. Hutchinson *et al.*, "Visual servo control, Part II: Advanced approaches," *IEEE Robotics and Automation Magazine*, vol. 14, no. 1, pp. 109–118, 2007.

[28] S. Onogi, T. Yoshida, Y. Sugano, T. Mochizuki, and K. Masuda, "Robotic ultrasound guidance by b-scan plane positioning control," *Procedia CIRP*, vol. 5, pp. 100–103, 2013.

[29] F. Aalamifar, R. Khurana, A. Cheng, R. H. Taylor, I. Iordachita, and E. M. Boctor, "Enabling technologies for robot assisted ultrasound tomography: system setup and calibration," in *SPIE Medical Imaging*. International Society for Optics and Photonics, 2014, pp. 90 401X–90 401X.

[30] S. Merouche, L. Allard, E. Montagnon, G. Soulez, P. Bigras, and G. Cloutier, "A robotic ultrasound scanner for automatic vessel tracking and three-dimensional reconstruction of b-mode images," *Ultrasonics, Ferroelectrics, and Frequency Control, IEEE Transactions on*, vol. 63, no. 1, pp. 35–46, 2016.

[31] B. Mustafa, A. Safwan, T. Ishii, Y. Matsunaga, R. Nakadate, H. Ishii, K. Ogawa, A. Saito, M. Sugawara, K. Niki *et al.*, "Development of robotic system for autonomous liver screening using ultrasound scanning device," in *Robotics and Biomimetics (ROBIO), 2013 IEEE International Conference on*. IEEE, 2013, pp. 804–809.

[32] G. Hirzinger, N. Sporer, A. Albu-Schaffer, M. Hahnle, R. Krenn, A. Pascucci, and M. Schedl, "Dlr's torque-controlled light weight robot iii-are we reaching the technological limits now?" in *Robotics and Automation, 2002. Proceedings. ICRA'02. IEEE International Conference on*, vol. 2. IEEE, 2002, pp. 1710–1716.

[33] R. Bischoff, J. Kurth, G. Schreiber, R. Koeppe, A. Albu-Schäffer, A. Beyer *et al.*, "The KUKA-DLR Lightweight Robot arm-a new reference platform for robotics research and manufacturing," in *Robotics (ISR)*. VDE, 2010, pp. 1–8.

[34] G. Schreiber, A. Stemmer, and R. Bischoff, "The Fast Research Interface for the KUKA Lightweight Robot," in *IEEE Workshop on Innovative Robot Control Architectures for Demanding (Research) Applications (ICRA 2010)*, 2010.

[35] L. C. Ebert, G. Hatch, G. Ampanozi, M. J. Thali, and S. Ross, "You Cant Touch This," *Surgical Innovation*, vol. 19, no. 3, pp. 301–307, 2012.

[36] C.-H. Hsieh, C.-H. Huang, and J.-D. Lee, "A Non-contact Image-to-Patient Registration Method Using Kinect Sensor and WAP-ICP," in *Software Engineering, AI, Networking and Parallel/Distributed Computing*. Springer, 2013, pp. 95–102.

[37] S. Billings, A. Kapoor, M. Keil, B. J. Wood, and E. Boctor, "A hybrid surface/image-based approach to facilitate ultrasound/CT registration," in *SPIE Medical Imaging*. International Society for Optics and Photonics, 2011, pp. 79 680V–79 680V.

[38] L. Mercier, T. Langø, F. Lindseth, and L. D. Collins, "A review of calibration techniques for freehand 3-d ultrasound systems," *Ultrasound in medicine & biology*, vol. 31, no. 2, pp. 143–165, 2005.

[39] S. Kahn, D. Haumann, and V. Willert, "Hand-eye calibration with a depth camera: 2D or 3D?" in *Computer Vision Theory and Applications (VISAPP), 2014 International Conference on*, vol. 3. IEEE, 2014, pp. 481–489.

[40] R. Y. Tsai and R. K. Lenz, "Real time versatile robotics hand/eye calibration using 3D machine vision," in *Robotics and Automation, 1988. Proceedings., 1988 IEEE International Conference on*. IEEE, 1988, pp. 554–561.

[41] S. Garrido-Jurado, R. Muñoz-Salinas, F. J. Madrid-Cuevas, and M. J. Marín-Jiménez, "Automatic generation and detection of highly reliable fiducial markers under occlusion," *Pattern Recognition*, vol. 47, no. 6, pp. 2280–2292, 2014.

[42] C. Herrera, J. Kannala, J. Heikkilä *et al.*, "Joint depth and color camera calibration with distortion correction," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 34, no. 10, pp. 2058–2064, 2012.

[43] R. C. Gonzalez, R. E. Woods, and S. L. Eddins, "Digital image processing using matlab," *Upper Saddle River, N. J: Pearson Prentice Hall*, 2004.

[44] J. Elseberg, D. Borrmann, and A. Nüchter, "One billion points in the cloud–an octree for efficient processing of 3D laser scans," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 76, pp. 76–88, 2013.

[45] J. K. Salisbury, "Active stiffness control of a manipulator in cartesian coordinates," in *Decision and Control including the Symposium on Adaptive Processes, 1980 19th IEEE Conference on*. IEEE, 1980, pp. 95–100.

[46] C. Hennersperger, A. Karamalis, and N. Navab, "Vascular 3D+ T freehand ultrasound using correlation of doppler and pulse-oximetry data," in *Information Processing in Computer-Assisted Interventions*. Springer, 2014, pp. 68–77.

[47] B. Fuerst, W. Wein, M. Müller, and N. Navab, "Automatic ultrasound–mri registration for neurosurgery using the 2D and 3D LC2 metric," *Medical image analysis*, 2014.

[48] M. R. Andersen, T. Jensen, P. Lisouski, A. K. Mortensen, M. K. Hansen, T. Gregersen, and P. Ahrendt, "Kinect depth sensor evaluation for computer vision applications," *Technical Report Electronics and Computer Engineering*, vol. 1, no. 6, 2015.